

Ghost artifact reduction for rendering 2.5D graphics

FIELD OF THE INVENTION

This invention relates to an image processing system for generating at least one output image associated with an output viewpoint from an input image associated with an input viewpoint through a depth-dependent transformation, the images being represented as
5 an input pixel array and an output pixel array, respectively.

This invention further relates to a method and a computer program product.

BACKGROUND OF THE INVENTION

The interest in providing a depth sensation when displaying an image on a 2D
10 display is growing rapidly, notably in virtual reality applications and computer games. Various forms of providing depth cues, such as influencing the brightness level or the size of an object, are known. Particularly, stereopsis, or stereoscopic vision, receives much attention as a technique for providing depth sensation. Stereopsis is evoked by presenting to a person two 2D images of the same scene as observed from two positions a little way apart. One of
15 the images is presented to the left eye, the other one is presented to the right eye. The two images are parallactically related. The term "parallax" refers to the apparent displacement or the difference in apparent direction of an object as seen from two different points not on a straight line with the object. Parallax allows a person to perceive the depth of objects in a scenery.

20 Many display techniques for displaying stereoscopic images are known. Using a time-parallel technique, both output images are presented simultaneously to one or two displays. For example, both images may be filtered with complementary colors and superimposed on one display. The observer wears glasses with filters that match the projection filters. Alternatively, both images may be displayed alongside on one display and
25 viewed using a viewer which directs each image to the correct eye. As another example, two displays may be used to present two differently polarized pictures, which are viewed through correspondingly polarized glasses. Alternatively, both images may be presented using a head-mounted device with separate displays for each eye. Also time-multiplexed techniques may be used, where the left and right images are alternatingly displayed on one display. As an

example, one image is written to the even scan lines of a monitor and the other image to the odd scan lines. A shutter system is used to occlude the left eye when the right-eye image is displayed and to occlude the right eye when the left-eye image is displayed. The shutter system may be mounted in glasses worn by the observer. Alternatively, a shutter with a 5 controllable polarizer is placed in front of the display and the observer wears a head-mounted device with polarized glasses.

A possible video format that is suited for rendering a 2D image from different viewpoints is a 2D video format enriched with depth information. Typically, the 2D input image is given in the form of an array of pixels. It may, for instance, have been obtained 10 using a camera or computer graphics. For each of the pixels of the input image additional depth information is available or, similarly, for each point in the scene, a distance is given of the point to the camera (or to another reference point, line or plane, such as a projection screen). Such a format is usually referred to as a 2.5D video format. The depth information allows modeling of the original image as a set of samples of a flat image, but also as a set of 15 samples of an image projected onto a terrain. Fig.2A shows a cross section of the sampled terrain. From the camera position a ray is cast to each pixel of the screen. The lengths of the arrows along the rays indicate the depth values of the pixel samples, also referred to as texels (texture elements). The 2.5D video format represents a subset of the full 3D model of the world. Rendering from other viewpoints can be accomplished by projecting the terrain onto 20 the image plane from the desired viewpoint. Based on the parallax when observing the same object from the visual point of the eyes and from the camera point, the pixels of the left eye image and the right eye image can be derived from the pixels of the input image by shifting the pixels. By choosing the horizontal axis of the coordinate system as being parallel to the line connecting both eyes, only a horizontal shift occurs. The amount of shift of a pixel 25 depends on the depth information. Fig.2B shows that after viewpoint transformation, the density of the projected input pixels is not uniform in the output domain. Hence, a resampling procedure is required. Existing video processing hardware/software designed for processing broadcast video signals is able to yield high-quality output images. Such hardware/software is able to scale an image to different display formats, such as 4:3 and 16:9 and perform 30 resampling and band limiting of the signal for such a conversion. The video processing usually occurs per horizontal scan line of the image.

SUMMARY OF THE INVENTION

It is an object of the invention to provide a video processing system and method capable of performing a viewpoint transformation without introducing visual artifacts.

To meet the object of the invention, an image processing system as described in the first paragraph is provided comprising: an input for receiving the input image and a hidden image, a video processor being operative to create output pixels of the output image and an output for providing the output image for subsequent rendering. The input image is a pre-filtered 2D representation of 3D objects as seen from the input viewpoint and comprises for each input pixel an associated input pixel value and an associated input pixel depth. The hidden image is another 2D representation of the 3D objects and comprises information, which information is at least partly occluded from the input viewpoint. The video processor is operative to create the output pixels by transforming each input pixel to a transformed input pixel, associated with the output viewpoint, as a function of the input pixel depth and creating the output image based on the transformed input pixels, using hidden image pixels for filling de-occluded areas and for at least one pixel position adjacent to the de-occluded areas for preventing ghost line artifacts, caused by transformation of the pre-filtered input image.

The input image may, for example, be obtained from the input viewpoint by a video camera and the hidden image may then be obtained from a viewpoint, different from the input viewpoint, by another camera. The input image and the hidden image may, for example, also be computer generated based on a 3D model of a particular scene. The hidden image may be a complete 2D representation of the 3D scene, as seen from a given viewpoint. The hidden image may also comprise only information, which information is occluded from the output viewpoint, and is required for only one or a few predetermined transformations.

De-occlusion may occur at depth discontinuities in the input image. Due to pre-filtering of the camera that records the original video, the pixels in the vicinity of the depth discontinuity contain color from both the foreground and the background objects in the input image. Therefore, simply filling the hole with de-occluded pixels from the hidden layer will result in a noticeable ghost line artifact. It is an advantage of the image processing system according to the invention that ghost line artifacts near the de-occluded hidden image pixels are prevented. This advantage is obtained by rendering a larger part of the hidden layer for replacing transformed input pixels adjacent to the de-occluded parts of the hidden layer. The hidden image pixels in the larger part of the hidden layer are inserted into the output image instead of the corresponding transformed input pixels.

In an embodiment of the invention the depth-dependent transformation is a transformation from the input viewpoint to a predetermined output viewpoint and the hidden image is associated with the output viewpoint. The predetermined distance may for example be a distance between a left and a right eye for providing a depth sensation. The input image 5 may be obtained by a first camera in a position of a left eye viewpoint and the hidden image by a second camera in a position of a right eye viewpoint. Not all information recorded by the second camera has to be stored in the hidden image. Only information in and adjacent to the areas that are de-occluded with a transformation from the left eye input viewpoint to the right eye output viewpoint has to be stored in the hidden image. After transformation of the input 10 image, this information is used to fill the de-occluded areas in the output image and to prevent ghost line artifacts near the de-occluded areas.

In another embodiment the hidden image is associated with the input viewpoint and the hidden image pixels are associated with a hidden image pixel value and a hidden image pixel depth, the video processor being operative to transform each hidden 15 image pixel to a transformed hidden image pixel, associated with the output viewpoint, as a function of the hidden image pixel depth, and to create the output image using transformed hidden image pixels for filling de-occluded areas and for at least one pixel position adjacent to the de-occluded areas. This embodiment is particularly useful for providing the possibility to perform various transformations over various distances. The input image may be obtained 20 from a first viewpoint by a first camera and the hidden image from a second viewpoint by a second camera. Preferably the information recorded by the second camera is then transformed to the first viewpoint. As a consequence the input image and the hidden image are thus associated with the same viewpoint. From this viewpoint both images may undergo 25 the same transformation to an arbitrary output viewpoint.

In a preferred embodiment the image processing system comprises pixel selection means for sequentially selecting input pixels per row. The pixel selecting means are arranged for selecting a hidden image pixel for pixel positions in a de-occluded area, for a first number of pixel positions before the de-occluded area and for a second number of pixels after the de-occluded area. For other pixel positions on the display line, transformed input 30 pixels are selected. The first and/or second number of pixel positions may depend on a width of a horizontal pre-filter, used during recording of the input image or may be determined based on an analysis of the input image.

Sequentially processing the pixels per row in a direction opposite a displacement from the input viewpoint to the output viewpoint makes it easy to detect

occlusion. This can be done by maintaining an x-coordinate extent indicating a furthest x-coordinate of the pixels and their footprints that have been transformed so far. If the transformation of the next footprint results in increasing the extent the transformed next footprint is at least partly not occluded. In this way it can be decided easily whether or not to exclude a pixel from the output image. Preferably, the x-axis is horizontal and the camera displacement also. If so desired, other directions are possible as well. The video processor is operative to determine that a hidden image pixel is de-occluded if a transformed input pixel increases the x-coordinate extent by more than a predetermined threshold. Preferably, if the x-coordinate of the input pixel after transformation is more than 1.5 pixels away from the last pixel position in the output image, the hole is too big and is filled with at least one pixel from the hidden layer. Since it is desired to use pixels from the hidden layer for pixel positions before the de-occluded area, a means is needed to detect the de-occlusion before the transformed input pixels of the input image are enclosed in the output image. Backtracking is usually not possible because color values are accumulated in the video filter. For early detection of the de-occluded area, the video processing may include maintaining a look ahead extent.

BRIEF DESCRIPTION OF THE DRAWINGS

In the drawings

Figure 1 shows a block diagram of a conventional system incorporating an image processing system;

Figures 2A-2B show a terrain viewed from different viewpoints;

Figures 3A-3C show a perspective projection;

Figures 4A-4B illustrate increased overlapping of objects when observed from different viewpoints;

Figures 5A-5B illustrate the appearance of a de-occluded area when objects are observed from different viewpoints;

Figure 6 shows a schematic representation of the occurrence of ghost artifacts;

Figure 7 shows a schematic representation of the preventing of ghost artifacts;

Figure 8 illustrates detection of occlusions, using an x-coordinate extent;

Figure 9 illustrates detection of de-occlusions, using a look ahead extent;

Figure 10 shows a block diagram of part of a video processing system according to the invention; and

Figure 11 shows a block diagram of a system in which the image processing system of the invention is advantageously used.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

To explain the system according to the invention, with reference to Figures 1 to 5, depth dependent transformation will be described. In this general description a pixel is regarded as a point (dimensionless). Figure 1 shows a block diagram of a conventional system in which the image processing system according to the invention may advantageously be used. The conventional system comprises a memory 100, such as a graphics memory, for storing a 2D input image 101. The input image 101 is comprised of an array of pixels divided in rows and columns. For each pixel a pixel value is given. Various ways of representing a pixel value are well known, such as an RGB (Red, Green, Blue) or YUV coding. The pixel value may be stored in full, e.g. using 16 or 24 bits per pixel. Alternatively, a Color Look-Up Table (CLUT) scheme may be used to code the pixel value using fewer bits, e.g. 8 bits. In addition to the pixel value, for each pixel a depth value is stored in the memory 100 as the input depth 102. The depth value may, for instance, be stored using 16 bits per pixel. If required, a separate memory may be used for storing the input depth 102. The input image 101 and the input depth 102 may be generated in any suitable way. As an example, two cameras located at different positions may be used, preferably each representing a different eye. From the two 2D-images obtained by the cameras, one image plus depth information can be formed. The depth information can then be supplied in addition to and, preferably, compatible with the conventional supply of only one 2D image, allowing the scene to be observed using either a conventional 2D display system or a stereoscopic display system. On a game computer or personal computer the input image 101 and the input depth 102 are usually generated by a 3D-rendering process which derives the information from a 3D model stored in a memory. Typically, the memory is part of the main memory of the computer. Communication means, such as a telecommunication means, audio/video broadcasting or a cable network, may be used to supply the input image 101 and the input depth 102 to the image processing system.

A processor 106 uses the input image 101 and the input depth 102 to generate at least one output image. In the example shown in Figure 1, a left image 107 and a right image 108 are generated. The left image 107 represents a 2D representation of the 3D scene as observed from an observation point coinciding with the left eye of an observer. Similarly, the right image 108 represents a 2D representation of the 3D scene as observed from an

observation point coinciding with the right eye of an observer. The processor 106 may build the output images in a memory 109, such as a graphics memory. The system according to the invention enables processing per pixel row, reducing the memory requirements. The memories can therefore be formed using line buffers for storing only one scan line of an image. As such, the image data can be processed as stream. If vertical filtering is applied, a few rows need to be stored. Usually a D/A converter 110 presents the output images on a suitable display 111, such as a stereoscopic display. In a broadcast receiver, the processor 106 may operate on the input image and input depth time-synchronized at pixel level to the supply of this information. The synchronization may be loosened by using a memory 103 for storing an input image 104 and input depth 105, being a respective copy of the input image 101 and input depth 102. The processor 106 then operates on the input image 104 and input depth 105, independent from the supply of the input image 101 and input depth 102. At suitable moments, e.g. when a complete new image has been created, the input image 101 and the input depth 102 are copied to the respective input image 104 and input depth 105. In a situation where the memory 100 and the memory 103 are physically combined in one memory block, the copying may be performed without physically copying the data, for instance by reassigning pointer registers. It will be appreciated that instead for maintaining a full image copy also a few pixels of a row or a few rows may be stored depending on the filter being applied.

The depth dependent transformation of the input image into the output image is associated with displacements of 3D objects relative to each other. The transformation described here is also referred as a parallactic transformation. The displacements occur as a consequence of, e.g., a change in the location of the observer relative to the scene, a change in the orientation of the observer relative to the scene, changing positions of the objects relative to each other due to their relative velocities, or a combination of these changes.

Figure 2A shows a terrain 203 seen from a camera viewpoint 201 and projected on an image plane 204. Normally, the image plane is formed by uniform pixels (with possibly a different density in x and y direction). In this example, for each pixel of the image plain the corresponding texel (texture element) of the terrain is determined (intersection of the ray from the camera to the pixel through the terrain). Depth may be represented in any suitable form (e.g. representing a length from the texel to the pixel, distance of texel to camera, distance of texel to projection plane, etc). The depth may also be coded in any suitable form, for instance, a 16-bit encoding on a linear scale may be used, where the depth of an object with minimum allowed depth is encoded as 0000H

(hexadecimal) and the depth of an object with maximum allowed depth is encoded as FFFFH. Persons skilled in the art will be able to select appropriate other representations, if so desired. Figure 2B illustrates what happens if the same terrain is viewed from a viewpoint 202. As will be clear, occlusion 206 (texel is no longer visible), minification 205 and 5 magnification 207 may occur. The same holds if the original object was not a terrain but true 3D objects.

Figure 3A shows a perspective projection. Shown is a 3D coordinate system with an x-axis 30, y-axis 31 and a z-axis 32. A 2D image is comprised of an array of discrete pixels arranged in rows and columns. Pixels in this context are the smallest entities 10 considered by the image processing according to the invention. Each respective pixel of a specific row in the image can assume only a respective one of a series of discrete positions. Each row of pixels in an image runs parallel to the x-axis 30, so that individual pixels in a row are discriminated on the basis of their respective x-coordinates. Each column of pixels runs parallel to the y-axis 31 that points in a direction perpendicular to the x-axis 30. The 15 depth of the scene is measured along the z-axis 32, which runs perpendicular to both the x-axis 30 and the y-axis 31. A respective z-value is assigned to each particular pixel in order to represent the depth of the scenery for that particular pixel. For explaining the parallactic shifting of pixels, the origin $O = (0, 0, 0)$ and the orientation of the coordinate system are chosen such that a 3D-object, represented by the point $P = (x, y, z)$ is observed from the 20 observation point $O_1 = (D, 0, 0)$. The plane $z = z_p$ is chosen to be the focus plane 33 (the plane on which the eyes of the observer focus). Usually, the plane of the display is chosen to coincide with the focus plane. The 2D image observed from the observation point O_1 is formed by a projection of the 3D objects on the focus plane, which is also referred to as the projection plane. The point P is projected onto $P_1 = (x_p, y_p, z_p)$.

25 Figure 3B shows a projection of Figure 3A onto the plane $y = 0$. P is projected onto $P' = (x, 0, z)$; $P'' = (D, 0, z)$ is a projection of P' on the plane $x = D$. From the triangle O_1, P', P'' it follows that: $z_p / z = (x_p - D) / (x - D)$, giving $x_p = D + (x - D) \cdot z_p / z$.

Figure 3C shows a projection of Figure 3A onto the plane $x = D$. P is projected onto $P' = (D, y, z)$; $P'' = (D, 0, z)$ is a projection of P' on the plane $y = 0$. From the triangle 30 O_1, P', P'' it follows that: $z_p / z = y_p / y$, giving $y_p = y \cdot z_p / z$.

This gives that $P_1 = (D + (x - D) \cdot z_p / z, y \cdot z_p / z, z_p)$. Similarly, $P_2 = (-D + (x + D) \cdot z_p / z, y \cdot z_p / z, z_p)$. Similar formulas are given in [IEEE Computer graphics & Applications, Tutorial: Time-Multiplexed Stereoscopic Computer Graphics, March 1992]. From the formulas, it follows that by choosing the x-axis to be parallel to the line through the

observation points, the y coordinates of P_1 and P_2 are the same. So, when deriving an image observed from O_2 from the image observed from O_1 , no vertical parallax occurs. In general, by choosing the x-axis in this way the calculation of the parallax is simplified. The pixels of the output image can be derived from the input image. Assuming that the input image

5 corresponds with the image as observed from O_1 and the output image corresponds with the image as observed from O_2 . Assuming further that for input pixel $p_i = (x_i, y_i)$ the pixel value is given, as well as the depth z_i of the 3D point $P(x, y_i, z_i)$ from which p_i is derived. The corresponding output pixel $p_o = (x_o, y_o)$, with $y_o = y_i$, is related to the same 3D point $P(x, y_i, z_i)$. This gives:

10
$$x_i = D + (x - D) \cdot z_p / z_i, \text{ and}$$

$$x_o = -D + (x + D) \cdot z_p / z_i.$$

This implies that x_o can be derived from x_i in the following way:

$$x_o = x_i - 2 \cdot D + 2 \cdot D \cdot z_p / z_i = x_i + 2 \cdot D \cdot (z_p / z_i - 1).$$

therefore, be derived from the input image by performing a horizontal shift d (only in the x-direction) of:

$$D = 2 \cdot D \cdot (z_p / z_i - 1). \quad (1)$$

From this formula it can be observed that the shift is proportional to the inverse of the depth. In the formula, $2 \cdot D$ corresponds to the offset between the observation points O_1 and O_2 .

20 Figure 4A shows a scene observed from an observation point O_3 , corresponding to the input image, wherein a nearby object 40 is partially overlapping a further object 41. For clarity, the y-coordinates are not shown in the figure. Due to the overlapping, the input image comprises all pixels of the nearby object 40 and only some pixels of the further object 41, as illustrated by the observed image 42. Figure 4B shows the
25 same scene observed from an observation point O_4 , which corresponds to the output image 43. As is clear from fig.4B the overlap has increased. If the output image 43 of Figure 4B is derived from the input image 42 of Figure 4A, this increased overlap will be reflected by at least one input pixel corresponding to the background object 41 partly or totally being occluded by an input pixel corresponding to the foreground object 40. In the areas of overlap
30 there are several candidates for the pixel values in the output image. The system according to the invention uses a reconstruction filter to determine for the output pixel locations (i.e. according to the grid of the output image) the input pixel(s) that after transformation contribute(s) to the output pixel location.

In contrast to different pixels of the input image being shifted to the same output pixel position, Figure 5 shows an example wherein a ‘hole’ occurs in the output image 53 due to the fact that no input pixel is present to fill a position in the output image 53. Figure 5A shows that in the input image 52 a background object 51 is partly obscured by a foreground object 50. In Figure 5B no overlap occurs between the objects 50 and 51 as can be observed from the output image 53, seen from an observation point O_6 .

With reference to Figure 6, the occurrence of ghost line artifacts when pixels from a hidden layer are included in an output image will be described. The image processing system receives an input image 601 and a hidden image 602 at the input. The input image 10 601 is associated with the input viewpoint and shows, for example, a dark foreground object 606 and light background object 607. Due to pre-filtering of the camera that records the original video, the pixels at the depth boundary 608 comprise some color of both, the dark foreground and the light background object. Depending on the chosen embodiment the hidden image 602 is either associated with the input viewpoint or the output viewpoint. 15 Transformation of the input image 601 to an output viewpoint may result in a de-occluded area 612 in the viewpoint transformed input image 603. This de-occluded area 612 is filled with pixels from the hidden image 602 and a filter is applied. It depends on the viewpoint the hidden image 602 is associated with, whether the hidden image 602 is to be transformed before the hidden image pixels can be used for filling the de-occluded area 612. Due to the 20 pre-filtering some pixels which before transformation were located at the original depth boundary in the input image will still comprise some color of both the foreground and the background object. The transformation results in a new depth boundary in the output image 605. As a consequence the output image 605 comprises two different depth boundaries and a ghost line artifact 610 appears.

25 With reference to Figure 7, the preventing of ghost line artifacts 610 as performed by the image processing system according to the invention will be described. In this system not only the pixel positions in the de-occluded area 609 are filled with pixels from the hidden image 602. Also one or more pixels adjacent to the de-occluded area are taken from the hidden image 602. As a consequence an enlarged de-occluded area 614 is 30 created and the pixels with color from both the foreground and the background are not part of the output image anymore. A filtering process is applied to the transformed pixels and an output image 613 without ghost artifacts is the result.

Figure 8 illustrates detection of occlusions, using an x-coordinate extent 80. Occlusions are for example detected as follows. The video processor 106 is operative to

sequentially process a scan line, i.e. a row with input pixels, in a direction opposite to a displacement from the input viewpoint v_0 81 to the output viewpoint v_1 82. The processing includes transforming input pixels to output pixels as a function of the associated input pixel depth. The projected pixels of pixels a 83, b 84 and c 85 are part of the input image

5 associated with the viewpoint v_0 81. With the transformation to the output image 86 associated with the viewpoint v_1 82 pixel a 83 gets occluded. The pixels b 84 and c 85 are not occluded and are mapped onto the output image 86. The processing may include maintaining an x-coordinate extent 80 that indicates for already processed input pixels with respect to a predetermined start position a furthest x-coordinate already occluded by at least one
10 transformed input pixel. Then, it can be concluded that if a pixel that is transformed does not lengthen the extent, it must be occluded by the previously processed pixels. In this event the pixel transformed from pixel a 83 does not lengthen the extent because it is occluded by the pixels transformed from pixel b 84. Therefore pixel a 83 is not included in the output image .
15 86. For a viewpoint transformations $v_1 < v_0$, the argument is analogous. In that event the scan line is traversed with increasing x-coordinate. Similarly, if a pixel that is transformed does lengthen the extent by more than a predetermined threshold, for example 1.5 pixels, part of the hidden image is de-occluded and at least one hidden image pixel is inserted in the output image 86.

In order to be able to efficiently fill in contributions from the hidden layer, it is
20 preferred to interleave processing of the base image layer, and the hidden layer. For both scans, the extent in the output scan line may be maintained. In this way, only a single scan over the base image scan line is performed interleaved with a single scan over the hidden image scan line. Specifying the pixels of the hidden layer in the coordinate frame of the base image allows for easy integration of the transformed input image and the transformed hidden
25 image.

In an embodiment of the image processing system according to the invention ghost line artifacts are detected as illustrated in Figure 9. Figure 9 shows the transformation of the input image shown in Figure 8, transformed to an output image associated with viewpoint v_2 ($v_2 > v_1$). In this embodiment of the image processing system, de-occlusions are
30 detected using a look ahead extent 91. The x-coordinate extent 80 maintains the input pixels which, after transformation, are fed to a filter that prefilters and samples the output scan line at screen resolution. The look ahead extent 91 which is ahead of the x-coordinate extent 80 detects possible areas of occlusion, minification, magnification and/or de-occlusion. If a pixel, after transformation, does lengthen the look ahead extent 91 by more than a

predetermined threshold, preferably 1.5 pixels, an area of de-occlusion is detected. For example after the viewpoint transformation from viewpoint v_0 to v_2 shown in Figure 9 a de-occluded area 92 occurs between the pixels b 84 and c 85. First the de-occlusion is detected by the look ahead extent 91. Later the x-coordinate extent 80 reaches the de-occluded area 5 and the de-occluded pixel d 93 from the hidden layer 94 is inserted in the output image 86. In order to prevent ghost line artifacts preference to the hidden layer 94 is given for pixel positions adjacent to the de-occluded hidden image pixels. Preferably one to four pixel positions before and/or after the de-occluded area are filled by pixels from the hidden image 94. In the filter color values of pixels are accumulated, which makes it extremely difficult to 10 replace already processed pixels. Therefore the detection of the de-occluded area has to take place before the input samples to be replaced are fed to the filter. Using a look ahead extent 91 enables early detection of a de-occluded area.

Figure 10 shows a schematic view of an embodiment 1000 of the image processing system according to the invention. The system uses four different extents. A base 15 x-coordinate extent 1001 and a base look ahead extent 1002 are used for transforming the input pixels and detecting de-occluded areas in the transformed input image 1010, respectively. A hidden x-coordinate extent 1003 and a hidden look ahead extent 1004 simultaneously do the same for the hidden image 1012. Preferably extra hidden layers are available and an x-coordinate extent and a look ahead extent are provided for each hidden 20 layer. Besides predetecting de-occluded areas the hidden look ahead extents may be used for determining if a hidden image pixel is available for inserting in or adjacent to a de-occluded area. When no hidden image pixel is available for filling a de-occluded area, the area may be filled using known methods such as linear interpolation. The video processor 106 is operative to detect for each input pixel 1011 of the input image 1010 how much the pixel lengthens the 25 base look ahead extent 1002 after transformation. If a pixel after transformation lengthens the look ahead extent by more than a predetermined threshold, for example 1.5 pixels, a de-occlusion is detected. Transformed pixels from the input image 1010 and the hidden image 1012 are fed to an input pixel selection unit 1005. The input pixel selection unit 1005 may be implemented as software or in hardware components in the video processor. The input pixel 30 selection unit 1005 selects, based on information from the look ahead extents, which transformed pixels will be part of the output image. The selected pixels are fed to a video filter block 1006 before being displayed on a suitable display. According to the invention, if a de-occluded area is detected by the look ahead extent 1002, near the de-occluded area preference is given to the transformed pixels of the hidden image 1012 when the input pixel

selection unit 1005 feeds pixels to the video filter block 1006. Hereby ghost artifacts near depth discontinuities are prevented. The pixel selection unit 1005 may be operative to determine if a hidden image pixel is available for inserting into the output image. If not, transformed input pixels known methods such as linear interpolation may be used for filling 5 the de-occluded area and the adjacent pixel positions.

Figure 11 shows a block diagram of a system in which the image processing system of the invention is advantageously used. The system shown in Figure 11 is an enhanced version of the conventional system shown in Figure 1. Additional to the input image 101 and the input depth 102 the memory 100 receives a hidden image 1101 and, 10 optionally, hidden image pixel depth values. If a memory 103 is used for storing a copy 104 of the input image 101 and a copy 104 of the input depth 102, this memory further stores copies 1104 and 1105 of the hidden image pixels 1101 and the hidden image pixel depths 1102. The processor 1106 is operative to create output pixels of the output image. The output pixels are created as described above. The processor 1106 may be operative to transform the 15 input image to an output image associated with a predetermined output viewpoint. In that event, the hidden image may not have to be transformed. In an embodiment of the system according to the invention the output viewpoint is chosen by an end user and both the input image pixels and the hidden image pixels are transformed to transformed pixels, associated with the output viewpoint.

It should be mentioned that the above mentioned embodiments illustrate rather than limit the invention, and that those skilled in the art will be able to design many alternative embodiments without departing from the scope of the appended claims. In the claims, any reference signs placed between parentheses shall not be construed as limiting the claim. Use of the verb "comprise" and its conjugations does not exclude the presence of 20 elements or steps other than those stated in a claim. The article "a" or "an" preceding an element does not exclude the presence of a plurality of such elements. The invention may be implemented by means of hardware comprising several distinct elements, and by means of a suitably programmed computer. In the device claim enumerating several means, several of 25 these means may be embodied by one and the same item of hardware. The mere fact that certain measures are recited in mutually different dependent claims does not indicate that a 30 combination of these measures cannot be used to advantage.